arXiv:1611.06981v1 [cs.CV] 21 Nov 2016

Multiple-View Spectral Clustering for Group-wise Functional Community Detection

Nathan D. Cahill[1,2⋆], Harmeet Singh[1], Chao Zhang[3,4], Daryl A. Corcoran[1,2], Alison M. Prengaman[1,2], Paul S. Wenger[2], John F. Hamilton[2], Peter Bajorski[2], and Andrew M. Michael[4]

[1] Image Computing & Analysis Laboratory (ICAL), RIT, Rochester, NY, USA
[2] School of Mathematical Sciences, RIT, Rochester, NY, USA
[3] Center for Imaging Science, RIT, Rochester, NY, USA
[4] Autism and Developmental Medicine Institute, Geisinger Health System, Danville, PA, USA

Abstract. Functional connectivity analysis yields powerful insights into our understanding of the human brain. Group-wise functional community detection aims to partition the brain into clusters, or communities, in which functional activity is inter-regionally correlated in a common manner across a group of subjects. In this article, we show how to use multiple-view spectral clustering to perform group-wise functional community detection. In a series of experiments on 291 subjects from the Human Connectome Project, we compare three versions of multiple-view spectral clustering: MVSC (uniform weights), MVSCW (weights based on subject-specific embedding quality), and AASC (weights optimized along with the embedding) with the competing technique of Joint Diagonalization of Laplacians (JDL). Results show that multiple-view spectral clustering not only yields group-wise functional communities that are more consistent than JDL when using randomly selected subsets of individual brains, but it is several orders of magnitude faster than JDL.

Keywords: spectral clustering, functional connectivity, fMRI

1 Introduction

In recent years, a variety of ideas have been proposed around the use of spectral clustering techniques to fuse group-wise information from structural and/or functional brain imagery. Van den Heuvel et al. [7] use a two-stage approach for identifying functional networks from resting-state fMRI (rsfMRI) that first clusters individual subjects and then partitions a graph constructed by connecting

⋆ Send correspondence to: nathan.cahill@rit.edu

individual clusters across the group of subjects. Chen et al. [4] fuse diffusion tensor imagery (DTI) and rsfMRI of multiple subjects using a co-training algorithm [9] that projects affinity matrices of individual subjects onto the eigenspaces of those of other subjects. While these approaches do yield useful results and insights, neither is stated as a well-posed optimization problem over the space of subject graphs, making them difficult to analyze from a mathematical perspective.

One recent approach that does formulate group-wise functional community detection as a well-posed optimization problem is described in Dodero et al. [5]: Joint Diagonalization of Laplacians (JDL) finds an optimal group-wise embedding of the graphs representing resting-state networks of each subject, and it subsequently applies a standard clustering technique to the optimal embedding in order to generate the functional communities. While JDL is a mathematically intriguing approach in its own right, the underlying algorithm [3] on which it relies requires computing an embedding in $\mathbb{R}^n$, where n is the number of graph vertices. Since the desired number of clusters, k, is typically far less than n, JDL may be very inefficient compared to algorithms that would only require embeddings in $\mathbb{R}^{k-1}$. Furthermore, the objective function minimized by JDL is not related to any specific graph partitioning cost, making it difficult to infer the true relative quality between two different clustering results.

In this article, we show how to use ideas from multiple-view spectral clustering [8,15] to design group-wise functional community detection algorithms that approximately minimize a well-defined cut cost and only require computing embeddings in $\mathbb{R}^{k-1}$. We then perform a series of experiments on a group of subjects from the Human Connectome Project [1] to show that multiple-view spectral clustering often yields more consistent clusters and is significantly more computationally efficient than JDL.

2 Preliminaries

Consider an undirected weighted graph $\mathcal{G} = (V, \mathcal{E})$ having n vertices that we wish to partition into k disjoint subgraphs $\mathcal{G}_i = (V_i, \mathcal{E}_i)$, $i = 1, 2, \ldots, k$, where $\bigcup_{i=1}^{k} V_i = V$. The graph $\mathcal{G}$ can be partitioned by removing the edges that connect each of the subgraphs to every other subgraph. A standard partitioning cost is the *multiclass normalized cut* cost [13], defined as a generalization of [11]:

$$\mathrm{NCut}_{\mathbf{W}}(V_1, \ldots, V_k) = \sum_{i=1}^{k} \frac{\mathrm{Cut}_{\mathbf{W}}(V_i, V \backslash V_i)}{\mathrm{Vol}_{\mathbf{W}}(V_i)} \ , \tag{1}$$

where $\mathbf{W}$ is the weighted adjacency matrix of $\mathcal{G}$, $\mathrm{Vol}_{\mathbf{W}}(V_i) = \sum_{v_j \in V_i} d_j$, and $d_j = \sum_{\ell} W_{j,\ell}$ is the degree of vertex v_j.

By defining an $n \times k$ indicator matrix $\mathbf{X}$ so that $X_{i,j} = 1$ if $v_i \in V_j$ and $X_{i,j} = 0$ otherwise, it is straightforward to see how the multiclass normalized cut cost can be expressed in terms of Rayleigh quotients. If $\mathbf{x}_i$ is the i^{th} column of $\mathbf{X}$, then $\mathrm{Vol}_{\mathbf{W}}(V_i)$ can be written in terms of the degree matrix $\mathbf{D} = \mathrm{diag}(\mathbf{d})$

as $\mathbf{x}_i^{\mathbf{T}}\mathbf{D}\mathbf{x}_i$, and the pairwise cut cost between V_i and $V\backslash V_i$ can be written as $\mathrm{Cut}_{\mathbf{W}}(V_i, V\backslash V_i) = \mathbf{x}_i^{\mathbf{T}}\mathbf{W}(\mathbf{1} - \mathbf{x}_i) = \mathbf{x}_i^{\mathbf{T}}\mathbf{d} - \mathbf{x}_i^{\mathbf{T}}\mathbf{W}\mathbf{x}_i = \mathbf{x}_i^{\mathbf{T}}\mathbf{D}\mathbf{x}_i - \mathbf{x}_i^{\mathbf{T}}\mathbf{W}\mathbf{x}_i = \mathbf{x}_i^{\mathbf{T}}(\mathbf{D} - \mathbf{W})\,\mathbf{x}_i$. This allows us to express (1) as:

$$\mathrm{NCut}_{\mathbf{W}}(V_1, \ldots, V_k) = \sum_{i=1}^{k} \frac{\mathbf{x}_i^{\mathbf{T}}(\mathbf{D} - \mathbf{W})\,\mathbf{x}_i}{\mathbf{x}_i^{\mathbf{T}}\mathbf{D}\mathbf{x}_i} = \mathrm{tr}\Big(\mathbf{X}^{\mathbf{T}}\mathbf{L}\mathbf{X}\big(\mathbf{X}^{\mathbf{T}}\mathbf{D}\mathbf{X}\big)^{-1}\Big) \quad , \quad (2)$$

where $\mathbf{L} = \mathbf{D} - \mathbf{W}$ is the Laplacian matrix of $\mathcal{G}$. Minimizing (2) is NP-hard; however, a fast approximate minimum can be found by relaxing the binary constraints on the entries of $\mathbf{X}$, minimizing the relaxed version of (2) by computing the generalized eigenvectors corresponding to the smallest $k-1$ nontrivial generalized eigenvalues of $\mathbf{L}\hat{\mathbf{x}} = \lambda\mathbf{D}\hat{\mathbf{x}}$, and discretizing the result by k-means clustering.

3 Multiple-View Spectral Clustering

Now consider a collection of m undirected weighted graphs $\mathcal{G}^{(\ell)} = \big(V, \mathcal{E}^{(\ell)}\big)$, $\ell = 1, \ldots, m$ that share a common set of vertices but different edges, and define $\mathbf{W}^{(\ell)}$ to be the weighted adjacency matrix for $\mathcal{G}^{(\ell)}$. Finding a common partitioning of all of the graphs, as in the single-view case, requires defining and optimizing a partitioning cost. A natural partitioning cost that generalizes (1) can be formulated in a manner that is equivalent to the spectral clustering technique of Zhou and Burges [15]:

$$\mathrm{NCut}_{\tilde{\mathbf{W}}}(V_1, \ldots, V_k) = \sum_{i=1}^{k} \frac{\mathrm{Cut}_{\tilde{\mathbf{W}}}(V_i, V\backslash V_i)}{\mathrm{Vol}_{\tilde{\mathbf{W}}}(V_i)} = \mathrm{tr}\left(\mathbf{X}^{\mathbf{T}}\tilde{\mathbf{L}}\mathbf{X}\Big(\mathbf{X}^{\mathbf{T}}\tilde{\mathbf{D}}\mathbf{X}\Big)^{-1}\right) \quad , \quad (3)$$

where $\tilde{\mathbf{W}} = \sum_{\ell=1}^{m} \alpha_\ell \mathbf{W}^{(\ell)}$, the α_ℓ's are nonnegative weights that sum to one, and $\tilde{\mathbf{D}}$ and $\tilde{\mathbf{L}}$ are the degree and Laplacian matrices, respectively, of the graph having weighted adjacency matrix $\tilde{\mathbf{W}}$. As described in [15], this partitioning cost arises when modeling a mixture of random walkers on undirected graphs. Its minimum yields a good cut on average across the collection of graphs, even though it may not be the best cut for the underlying individual graphs in the collection.

In practice, minimizing (3) is NP-hard, but as with (2), a relaxed version of (3) can be solved by identifying the generalized eigenvectors corresponding to the smallest $k-1$ nontrivial generalized eigenvalues of $\tilde{\mathbf{L}}\hat{\mathbf{x}} = \lambda\tilde{\mathbf{D}}\hat{\mathbf{x}}$, and the result can subsequently be discretized by k-means clustering.

Choice of Weights: The choice of $\alpha_1, \ldots, \alpha_m$ is a hyperparameter that must be determined by the user. If uniform weights are chosen ($\alpha_j = m^{-1}$, $j = 1, \ldots, m$), we refer to the multiple-view spectral clustering algorithm as MVSC. Another option is to include weight selection in the optimization itself, for example, via the line search technique described in Huang et al. [8]. Huang et al. refer to this approach as Affinity Aggregation for Spectral Clustering (AASC).

We propose a third option, motivated from the idea that weights should be chosen for each individual graph in proportion to the quality of its embedding. Zhang and Jordan [14] show that the minimum normalized cut cost for $\mathcal{G}^{(\ell)}$ is equal to the sum of the smallest $k-1$ nontrivial generalized eigenvalues satisfying $\mathbf{L}^{(\ell)}\mathbf{x} = \lambda \mathbf{D}^{(\ell)}\mathbf{x}$. If we call these eigenvalues $\lambda_1^{(\ell)}$, $\lambda_2^{(\ell)}$, ..., $\lambda_{k-1}^{(\ell)}$, then the choice of weights:

$$\alpha_j = \frac{\left[\sum_{i=1}^{k-1} \lambda_i^{(j)}\right]^{-1}}{\sum_{\ell=1}^{m}\left[\sum_{i=1}^{k-1} \lambda_i^{(\ell)}\right]^{-1}} \ , \quad j = 1, \ldots, m \ , \tag{4}$$

will ensure that individual graphs with high quality partitioning cost will be more heavily weighted in the computation of the group-average Laplacian than individual graphs with lower quality partitioning cost. In the sequel, we will refer to multiple-view spectral clustering with weights given by (4) as MVSCW.

Relationship to Joint Diagonalization of Laplacians: The recent group-wise functional community detection technique proposed by Dodero et al. [5] identifies an orthogonal matrix $\mathbf{Q}$ that approximately jointly diagonalizes all of the normalized graph Laplacian matrices by minimizing:

$$\varepsilon(\mathbf{Q}) = \sum_{\ell=1}^{m} \mathit{off}\Big(\mathbf{Q}^{\mathbf{T}}\mathbf{D}^{(\ell)^{-1/2}}\mathbf{L}^{(\ell)}\mathbf{D}^{(\ell)^{-1/2}}\mathbf{Q}\Big) \ , \tag{5}$$

subject to the orthogonality constraint $\mathbf{Q}^{\mathbf{T}}\mathbf{Q} = \mathbf{I}$, where $\mathit{off}(\mathbf{A}) = \sum_{i \neq j} |a_{i,j}|^2$.

While this approach does enable the use of a simultaneous diagonalization algorithm that has a theoretical guarantee of convergence [3], the cost function (5) does not directly quantify the cost of partitioning the graph vertices. To see why this is true, consider the partitioned matrix $\mathbf{Q} = [\mathbf{Q}_1\ \mathbf{Q}_2]$, where $\mathbf{Q}_1$ is $n \times k$ and $\mathbf{Q}_2$ is $n \times (n-k)$. Suppose the rows of $\mathbf{Q}_1$ form an embedding of the graph vertices that is rounded to form the cluster indicator matrix $\mathbf{X}$. Then, in order for (5) to be thought of as approximately modeling some graph partitioning cost, $\varepsilon(\mathbf{Q})$ should be invariant with respect to any transformation of the form $\mathbf{Q}_2 \leftarrow \mathbf{Q}_2\mathbf{U}$, where $\mathbf{U}$ is any orthogonal $(n-k) \times (n-k)$ matrix. However, (5) does not exhibit this type of invariance.

4 Experiments

In this section, we compare three different multiple-view spectral clustering algorithms for group-wise community detection (MVSC, MVSCW, AASC) with Joint Diagonalization of Laplacians (JDL). We use each of the four algorithms to compute embeddings, and then we generate labellings via k-means clustering. Clustering is performed 100 times with different random seeds; the mode of each vertex is selected as the label for that vertex. Since labellings are ambiguous up to permutation, every comparison between two or more labellings includes

a step of identifying the permutation that maximizes the relative overlap (Dice coefficient) between labellings.

Data: To compare algorithms, we utilize rsfMRI data comprising one run from each of 291 female participants from the S500 release of the Human Connectome Project (HCP) [1]. All images were collected on a 3T Siemens Skyra scanner with a 32-channel receive head coil. T2* weighted functional images were acquired using a gradient-echo EPI sequence with TE = 33.1 ms, TR = 0.72 s, flip angle = 52°, slice thickness = 2 mm, field of view = 208 × 180 mm, matrix size = 104 × 90, voxel size = $2 \times 2 \times 2$ mm^3. Scan duration was approximately 15 minutes, yielding $\sim$ 1200 volumes. HCP provides preprocessed rsfMRI data [6]; we select the data that has been subsequently post-processed according to the FIX protocol [10] for ICA-based denoising.

Each FIX-preprocessed rsfMRI is parcellated according to the Automatic Anatomical Labeling (AAL) atlas [12], which divides the brain into 90 cortical/subcortical regions and 26 cerebellar/vermis regions. Time series from voxels within each region are averaged to form a representative time series for each region. The weighted adjacency matrices $\mathbf{W}^{(\ell)}$ are constructed from Fisher z-transformed Pearson correlation coefficients between the representative time series for each region. Negative weights are zeroed, ensuring all graph Laplacian matrices remain positive semi-definite.

Number of Clusters: The true underlying number of functional clusters in the brain is unknown. For these experiments, we choose the number of clusters based on the presence of an eigenvalue gap in the group-averaged Laplacian when using the full set of 291 female brains. As shown in Figure 1, when the MVSC and JDL algorithms are used, there is a large gap between the 4^{th} and 5^{th} eigenvalues and a smaller gap between the 7^{th} and 8^{th} eigenvalues. For this reason, we compute two versions of group-averaged Laplacians for use in the MVSCW algorithm: one with the choice $k = 5$, and the other with $k = 8$. (Note that the weights in MVSCW depend on the number of desired clusters.) For both of these cases, MVSCW exhibits similar gaps between the 4^{th} and 5^{th} and between the 7^{th} and 8^{th} eigenvalues. When the AASC algorithms are used, there is a gap between the 7^{th} and 8^{th} eigenvalues, but no discernible gap between the 4^{th} and 5^{th} eigenvalues. The same eigenvalue gaps are apparent when analyzing a similar set of 203 male brains. For these reasons, we perform subsequent experiments for both $k = 5$ and $k = 8$ clusters. The choice of eight clusters is consistent with number of clusters found by the co-training algorithm of Chen et al. [4].

Clustering Results: Figures 2–3 illustrate the group-wise functional communities that result from applying MVSC/MVSCW/AASC/JDL + k-means clustering to the 291 female brains for the cases of $k = 5$ and $k = 8$ clusters. By inspecting the regions of the brain that are clustered into similar communities, we see that for the $k = 5$ case, communities roughly correspond to the sensorimotor network, the default mode network, the visual network, and the

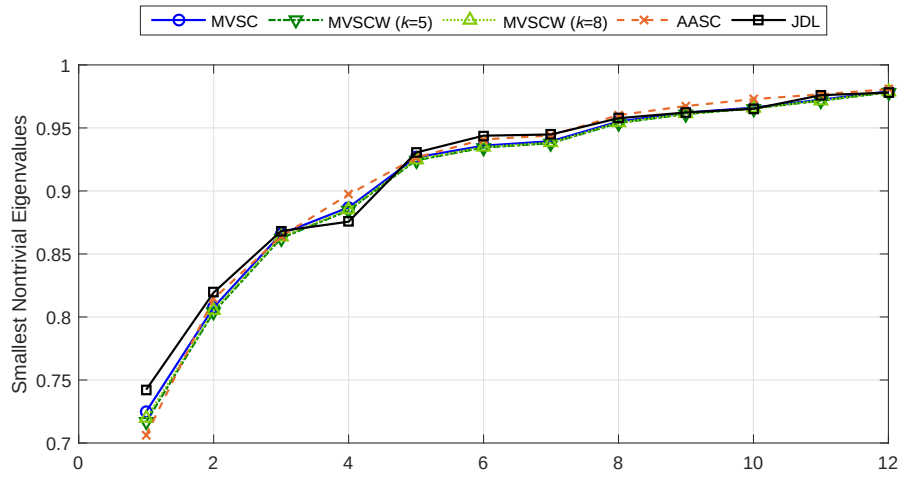

Fig. 1. Smallest nontrivial generalized eigenvalues computed on 291 female brains using MVSC, MVSCW, AASC, and JDL algorithms. A large gap is apparent between the 4^{th} and 5^{th} eigenvalues for MVSC, MVSCW, and JDL, and a smaller gap is apparent between the 7^{th} and 8^{th} eigenvalues for all algorithms.

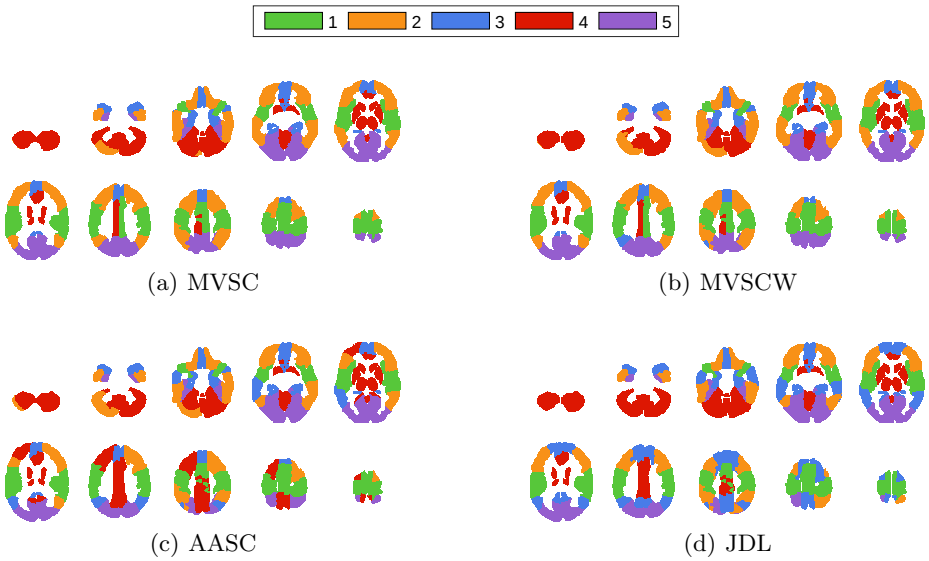

(a) MVSC (b) MVSCW

(c) AASC (d) JDL

Fig. 2. Group-wise clusters ($k = 5$). Clusters roughly correspond to: 1) Sensorimotor network, 2) Default mode network (MVSC/MVSCW), 3) Orbitofrontal cortex network (MVSC/MVSCW), 4) Basal ganglia + cerebellum, and 5) Visual network. For AASC/JDL, clusters 2 and 3 roughly correspond to the union of the orbitofrontal cortex and default mode networks.

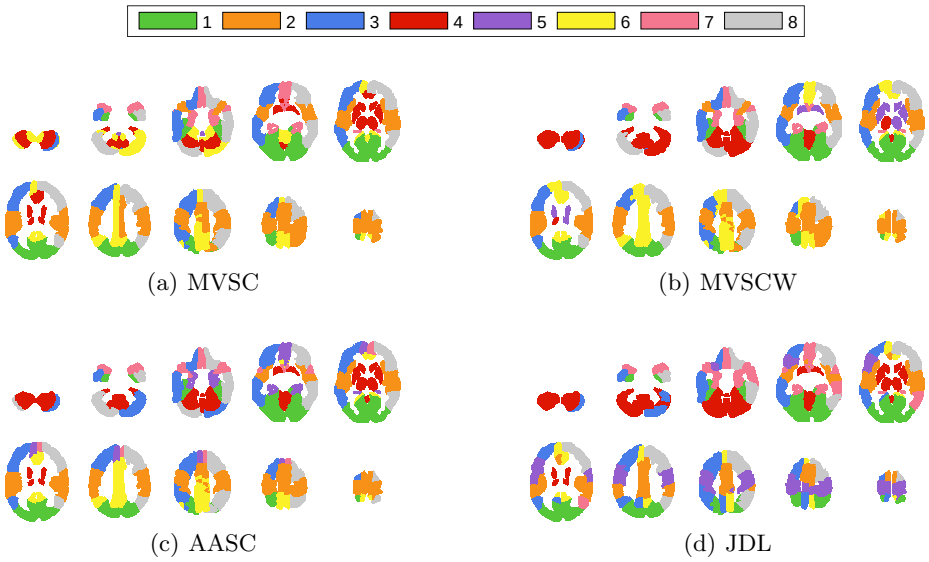

(a) MVSC (b) MVSCW

(c) AASC (d) JDL

Fig. 3. Group-wise clusters ($k = 8$). Some of the identified clusters roughly correspond to: 1) Visual network, 2) Sensorimotor network, 3) Left frontoparietal network, 4) Basal ganglia + cerebellum, 7) Orbitofrontal cortex network, and 8) Right frontoparietal network.

combined basal ganglia and cerebellar structures. However, for AASC and JDL, the default mode and orbitofrontal cortex networks are not as nicely separated as they are by MVSC and MVSCW. For the $k = 8$ case, communities emerge that roughly correspond to the sensorimotor, visual, and left and right frontoparietal networks, and the combined basal ganglia and cerebellar structures. The default mode network does not appear to emerge as a unique community for any of the algorithms for the $k = 8$ case, and some of the communities (5,6) do not appear to correspond to recognizable networks.

Consistency: Even though it is impossible to intuit the "correct" clustering of this data into functional communities, it is possible to assess how consistently each algorithm yields similar communities. To compare the consistency of the MSVC, MVSCW, AASC, and JDL algorithms (with k-means clustering applied to each embedding), we carry out the following experiment: For group sizes of $\gamma = 4$, 8, 16, 32, 64, and 128, we randomly select two non-overlapping subsets, each containing γ of the 291 brains. For each subset in a pair, we compute functional communities using MSVC, MVSCW, AASC, and JDL embeddings followed by k-means clustering with 100 random seeds. Then, for each algorithm, we compute the Dice coefficient (relative overlap) between the labellings generated from the two subsets in the pair (after identifying the permutation of one of the labellings that gives maximal overlap with the other labeling). This

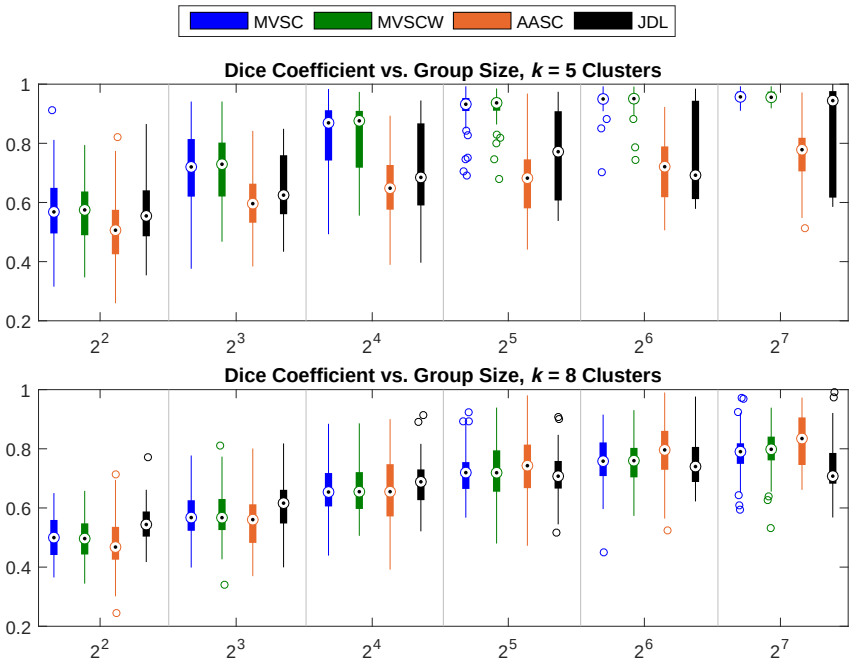

Fig. 4. Box plots of Dice coefficients (relative overlap in labellings) between pairs of labellings generated from MVSC/MVSCW/AASC/JDL + k-means clustering using groups of brains (group sizes 4, 8, ..., 128) randomly drawn from the 291 female brains.

entire process is performed 100 times, both for $k = 5$ and $k = 8$ clusters. Box plots of the resulting Dice coefficients are shown in Figure 4.

From the results of these experiments, we can see that for $k = 5$ clusters, MVSC and MVSCW yield more consistent labellings than AASC and JDL across all group sizes. For $k = 8$ clusters, JDL is more consistent than MVSC, MVSCW, and AASC for group sizes up to 16, and AASC is the least consistent for these small group sizes. For group sizes greater than 16, however, JDL becomes the least consistent and AASC the most consistent. For both $k = 5$ and $k = 8$, MVSC and MVSCW seem to exhibit similar consistency.

Timing: We would expect that of the four algorithms, MVSC, MVSCW, and AASC should be faster than JDL due to the fact that they only need to compute k eigenvalues, whereas JDL must compute all n eigenvalues. MVSC should be faster than MVSCW due to the weight computation in MVSCW, and it should be faster than AASC, since AASC successively solves the same subproblem that MVSC solves once. We implemented the MVSC and MVSCW algorithms in MATLAB R2015b, and we used the MATLAB implementations of AASC and JDL that are provided by their respective authors. When running

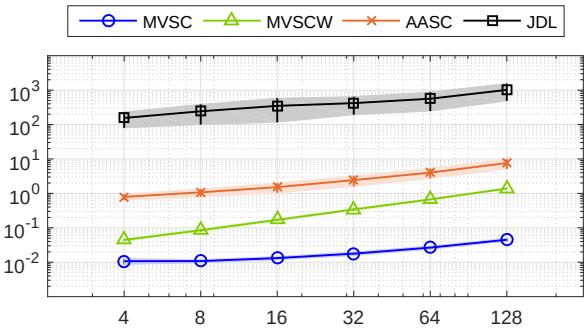

Fig. 5. Computation time (in seconds) required to compute MVSC, MVSCW, AASC, and JDL embeddings with $k = 8$ for various group sizes. Shaded regions indicate $\pm$ one standard deviation of the mean timing result.

the cluster consistency experiment on our quad-core desktop, we captured timing results for each algorithm, and we show the results in Figure 4. Timing results reflect the amount of time required to compute the group-wise embedding that is subsequently input into the k-means clustering algorithm, given the set of individual graph weighted adjacency matrices (we chose $k = 8$ clusters for this experiment). As seen in the figure, MVSC is four orders of magnitude faster than JDL, two orders of magnitude faster than AASC, and one order of magnitude faster than MVSCW across all group sizes, and each algorithm appears to grow linearly in complexity with respect to the number m of component graphs in the group.

5 Conclusion

When modeled as a graph partitioning problem, group-wise functional community detection from brain fMRI can be carried out by performing multiple-view spectral clustering algorithms. Experiments on a set of 291 female brains show that MVSC/MVSCW/AASC yield functional communities that roughly correspond to a variety of known functional networks in the brain. In addition, when compared to the previously proposed joint diagonalization of Laplacians (JDL) technique, multiple-view spectral clustering can yield more consistent results with much faster computation times.

Appendix

Prototype implementations of the MVSC and MVSCW algorithms, as well as wrappers around the original authors' implementations of AASC and JDL, are available for download at MATLAB Central [2] under File ID #58753.

References

1. Human Connectome Project. https://db.humanconnectome.org/
2. MATLAB Central. http://www.mathworks.com/matlabcentral/
3. Cardoso, J.F., Souloumiac, A.: Jacobi angles for simultaneous diagonalization. SIAM Journal on Matrix Analysis and Applications 17(1), 161–164 (1996)
4. Chen, H., Li, K., Zhu, D., Jiang, X., Yuan, Y., Lv, P., Zhang, T., Guo, L., Shen, D., Liu, T.: Inferring group-wise consistent multimodal brain networks via multi-view spectral clustering. IEEE Transactions on Medical Imaging 32(9), 1576–1586 (2013)
5. Dodero, L., Gozzi, A., Liska, A., Murino, V., Sona, D.: Group-wise functional community detection through joint Laplacian diagonalization. In: Proc. Medical Image Computing and Computer-Assisted Intervention (MICCAI). pp. 708–715. Springer (September 2014)
6. Glasser, M.F., et al.: The minimal preprocessing pipelines for the Human Connectome Project. Neuroimage 80, 105–124 (2013)
7. van den Heuvel, M., Mandl, R., Hulshoff Pol, H.: Normalized cut group clustering of resting-state fMRI data. PLoS ONE 3(4), e2001 (2008)
8. Huang, H.C., Chuang, Y.Y., Chen, C.S.: Affinity aggregation for spectral clustering. In: Proc. Computer Vision and Pattern Recognition (CVPR). pp. 773–780. IEEE (2012)
9. Kumar, A., Daumé, H.: A co-training approach for multi-view spectral clustering. In: Proc. International Conference on Machine Learning (ICML). pp. 393–400. ACM (2011)
10. Salimi-Khorshidi, G., Douaud, G., Beckmann, C.F., Glasser, M.F., Griffanti, L., Smith, S.M.: Automatic denoising of functional MRI data: combining independent component analysis and hierarchical fusion of classifiers. Neuroimage 90, 449–468 (2014)
11. Shi, J., Malik, J.: Normalized cuts and image segmentation. IEEE Transactions on Pattern Analysis and Machine Intelligence 22(8), 888–905 (2000)
12. Tzourio-Mazoyer, N., Landeau, B., Papathanassiou, D., Crivello, F., Etard, O., Delcroix, N., Mazoyer, B., Joliot, M.: Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. Neuroimage 15(1), 273–289 (2002)
13. Yu, S.X., Shi, J.: Multiclass spectral clustering. In: Proc. International Conference on Computer Vision (ICCV). pp. 313–319. IEEE (2003)
14. Zhang, Z., Jordan, M.I.: Multiway spectral clustering: A margin-based perspective. Statistical Science 23(3), 383–403 (2008)
15. Zhou, D., Burges, C.J.C.: Spectral clustering and transductive learning with multiple views. In: Proc. International Conference on Machine Learning (ICML). pp. 1159–1166. ACM (2007)